

Fuzzy Bayesian Learning

Indranil Pan and Dirk Bester

Abstract—In this paper we propose a novel approach for learning from data using rule based fuzzy inference systems where the model parameters are estimated using Bayesian inference and Markov Chain Monte Carlo (MCMC) techniques. We show the applicability of the method for regression and classification tasks using synthetic data-sets and also a real world example in the financial services industry. Then we demonstrate how the method can be extended for knowledge extraction to select the individual rules in a Bayesian way which best explains the given data. Finally we discuss the advantages and pitfalls of using this method over state-of-the-art techniques and highlight the specific class of problems where this would be useful.

Index Terms—fuzzy logic; mamdani method; machine learning; MCMC.

I. INTRODUCTION

PROBABILITY theory and fuzzy logic have been shown to be complementary [1] and various works have looked at the symbiotic integration of these two paradigms [2], [3] including the recently introduced concept of Z-numbers [4]. Historically fuzzy logic has been applied to problems involving imprecision in linguistic variables, while probability theory has been used for quantifying uncertainty in a wide range of disciplines. Various generalisations and extensions of fuzzy sets have been proposed to incorporate uncertainty and vagueness which arise from multiple sources. For example, the type-2 fuzzy [5], [6] sets and type-n fuzzy sets [5] can include uncertainty while defining the membership functions themselves. Intuitionistic fuzzy sets [7] additionally introduce the degree of non-membership of an element to take into account that there might be some hesitation degree and the degree of membership and non-membership of an element might not always add to one. Non-stationary fuzzy sets [8] can model variation of opinion over time by defining a collection of type 1 fuzzy sets and an explicit relationship between them. Fuzzy multi-sets [9] generalise crisp sets where multiple occurrences of an element are permitted. Hesitant fuzzy sets [10] have been proposed from the motivation that the problem of assigning a degree of membership to an element is not because of a margin of error (like Atanassov’s intuitionistic fuzzy sets) or a possibility distribution on possibility values (e.g. type 2 fuzzy sets), but because there is a set of possible values [10]. Formally these can be viewed as fuzzy multi-sets but with a different interpretation. Apart from these quantitative tools, other approaches have been used to address the vagueness of the qualitative meanings used by experts like fuzzy linguistic approaches [11] and various linguistic models have been proposed to address this [12]–[14].

Bayesian inference, on the other hand, provides a method of incorporating the subjective beliefs of the experts about the value of a model parameter, using prior probability distributions, through expert elicitation methods [15]. However, expert knowledge does not always come in the form of precise probability distributions and often have vague and ambiguous statements embodying some form of knowledge [16]. In such cases, both uncertainty and imprecision exist simultaneously [17]. The data is said to be uncertain if the degree of confidence in what is represented by the data is less than one [17], [18]. On the other hand data is imprecise if the implied attribute is not singular but belongs to an ill defined set [17], [18]. These two things cannot be discerned if they are represented by a single probability distribution and the contribution of each of these in the final posterior distribution cannot be tracked [19].

However, fewer works have looked at integration of Bayesian methods with fuzzy logic [20], [21]. For example, [22] attempts to generalise Bayesian methods for samples of fuzzy data and for prior distributions with imprecise parameters. Recent works have also looked at extension of these works for possibilistic Bayesian inference [23]. However, in all the above-mentioned literatures, most of the work is on representing the data in terms of fuzzy numbers and optimising model parameters based on that. None of the literatures specifically leverage on one of the fundamental strengths of the fuzzy inference system, i.e. expressing the expert opinion in terms of rule bases. This is specifically what this paper addresses and the idea stems from the fact that often the experts not only have an opinion on the uncertainty estimates of the model covariates, but also have an intuition of the inter-dependence between the various model covariates on the final output. None of the above-mentioned methods take this explicitly into account and hence do not fully exploit the embodied knowledge provided by the experts.

Another line of research has been the use of the belief rule base (BRB) inference methodology as proposed in [24] which uses an expert rule base and assigns a belief degree to each rule, individual rule weights and attribute weights. This treatment is not fully Bayesian and including these additional parameters for characterising these rule bases requires multi-objective optimisation methods to give acceptable levels of performance [25]. Having to estimate more number of parameters from a small data-set is inherently difficult and would have large confidence intervals (using frequentist methods) or large highest density intervals (using Bayesian methods). This issue is addressed in our methodology, since it does not need to have too many additional parameters over and above that required by the fuzzy rule base itself. This also helps in reducing the explosion of parameters if higher number of rule bases or fuzzy sets are used. Also the outputs in these

I. Pan was with Imperial College London, UK, and now works at Sciemus Ltd, London UK. e-mail: i.pan11@imperial.ac.uk, ipan@sciemus.com

D. Bester was with University of Oxford, UK, and now works at Sciemus Ltd, London UK. e-mail: dbester@sciemus.com

works is a belief distribution over the fuzzy variables like low, medium and high, instead of a distribution on the crisp output of real numbers, as proposed later in this work. This is important especially if such models are developed as a part of a larger decision making framework which relies on probability distributions as outputs from the sub-models. A Bayesian extension of this framework is proposed in [26] to update the belief degrees based on a recursive online estimation. This also addresses a different problem and has the same issues as the BRB methodology mentioned earlier.

Adaptive fuzzy systems have been used to approximate non-conjugate priors, likelihoods and approximate hyper-priors in Bayesian inference, which helps in allowing users to extend the inference for a wider range of prior and likelihood probability density functions [27]. In our proposed method, we use Markov Chain Monte Carlo (MCMC) [28] to numerically approximate the posterior distribution and hence do not have the issues relating to conjugacy. Consequently, our method can work on any likelihood or prior distribution.

The motivation of the present work stems from the need to incorporate linguistic expert opinion within the standard Bayesian framework, so that the end-to-end overall system works on a set of crisp input-output data, which might have probabilistic uncertainties associated with them. It is easier for experts with domain knowledge to come up with a linguistic rule base, but to get good prediction accuracy on any data-set, the shape of the membership functions need to be appropriately chosen. The novelty of our work is that we convert this problem into a Bayesian parameter estimation problem where the shape of the membership functions are estimated through an MCMC algorithm for a given data-set. We also extend this to Bayesian inference on the individual rules themselves to find out which of the rules are more probable, given the data.

This approach leverages on the strengths of both the disciplines – statistics and computational intelligence to come up with a practical solution methodology. In traditional statistics, the emphasis is on parameter estimation for a given model structure (which is chosen by the user) from a given set of data-points. In soft computing, the emphasis is on mathematically laying down the imprecision in the spoken language and expert opinions in the form of rule bases and does not need data to generate membership functions at the outset (data might be used later for calibration of the membership functions to improve prediction accuracies). Both of these advantages are synergistically used in the present work and the subsequent section develops this proposition in more detail.

II. PROPOSED METHODOLOGY AND DEMONSTRATIVE EXAMPLES

All empirical modelling approaches which learn from data, like linear and non-linear regression, generalised linear models, neural networks, or kernel based methods like Gaussian processes, support vector machines etc. can be interpreted using a generic Bayesian framework as an inference of the underlying non-linear model function $g(\mathbf{x}; \theta)$ (with parameters θ), given the data. For the set of input vectors $\mathbf{X}_N = \mathbf{x}_i \forall i \in \{1, 2, \dots, N\}$ and the set of outputs $\mathbf{Y}_N = \mathbf{y}_i \forall i \in \{1, 2, \dots, N\}$

the inference of the function $g(\mathbf{x}; \theta)$ can be mathematically represented by the posterior probability distribution [29]

$$p(g(\mathbf{x}; \theta) | \mathbf{Y}_N, \mathbf{X}_N) = \frac{p(\mathbf{Y}_N | g(\mathbf{x}; \theta), \mathbf{X}_N) p(g(\mathbf{x}; \theta))}{p(\mathbf{Y}_N | \mathbf{X}_N)} \quad (1)$$

where, $p(\mathbf{Y}_N | g(\mathbf{x}; \theta), \mathbf{X}_N)$ is the likelihood function which gives the probability of the output values given the function $g(\mathbf{x}; \theta)$ and the term $p(g(\mathbf{x}; \theta))$ is the prior distribution of the functions and its corresponding parameters assumed by the model. As with most other modelling approaches, the above equation assumes that the covariates are independent of the model parameters and the functional form of $g(\mathbf{x}; \theta)$, i.e., $p(\mathbf{X}_N | g(\mathbf{x}; \theta)) = p(\mathbf{X}_N)$, but the assumption breaks if there is selection bias [15]. Any choice on the parametric structure of the model $g(\mathbf{x}; \theta)$, implicitly specifies $p(g(\mathbf{x}; \theta))$ in terms of the non-linear relationships among the covariates, smoothness, continuity etc.

For parametric models like linear regression where the structure of the model $g(\mathbf{x}; \theta)$ is known, Eq. (1) essentially transforms into an inference only on the model parameters θ given by

$$p(\theta | \mathbf{Y}_N, \mathbf{X}_N) = \frac{p(\mathbf{Y}_N | \theta, \mathbf{X}_N) p(\theta)}{p(\mathbf{Y}_N | \mathbf{X}_N)} \quad (2)$$

For non-parametric approaches, like Gaussian processes, the prior $p(g(\mathbf{x}; \theta))$ can be placed directly on the function space without any requirement of explicit parameterisation of the function $g(\mathbf{x}; \theta)$. In our case, the function $g(\mathbf{x}; \theta)$ is a fuzzy inference system given by the rule base Eq. (3)

$$\beta_k R_k : \text{if } A_1^k \oplus A_2^k \oplus \dots \oplus A_{T_k}^k \text{ then } C_k \quad (3)$$

where β_k is a dichotomous variable indicating the inclusion of the k^{th} rule in the fuzzy inference system, $A_i^k \forall i \in \{1, 2, \dots, T_k\}$ is a referential value of the i^{th} antecedent attribute in the k^{th} rule, T_k is the number of antecedent attributes used in the k^{th} rule, C_k is the consequent in the k^{th} rule, $\oplus \in \{\vee, \wedge\}$ represents the set of connectives (OR, AND operations) in the rules. A referential value is a member of a referential set with the universe of discourse U_i , which are meaningful representations for specifying an attribute using subjective linguistic terms which are easily understood by an expert. For example, an attribute representing maintenance level might consist of linguistic terms like ‘poor’, ‘medium’ and ‘good’. Each of these referential values are characterised by membership functions $f_i(A_i, U_i, \phi_i)$ or $f_i(C_i, U_i, \phi_i)$ and for the inference of the function $g(\mathbf{x}; \theta)$ in Eq. (1), the set of parameters to be estimated is θ , which contains all ϕ_i and β_k

Remark 1: Since the parameter estimation problem does not impose any specific structure on Eq. (3), the underlying fuzzy rule bases or sets can be defined to be of other types like fuzzy type II, intuitionistic fuzzy sets, hesitant fuzzy sets etc.

Remark 2: Including β_k in Eq. (3) for selecting each rule and estimating these co-efficients within a Bayesian framework, essentially results in knowledge extraction from the data in the form of rules, i.e., the expert has some prior opinion on specific set of rules and the Bayesian inference

engine combines the data with this to show which of the ones are more likely given the data.

A. Regression on synthetic data-sets

1) *Bayesian estimation of fuzzy membership function parameters:* Consider the synthetic example of predicting downtime of an engineering process system based on the variables location risk and maintenance level. For the referential sets *loc_risk*, *maintenance*, *downtime* the referential values are $\{LO, MED, HI\}$, $\{POOR, AVG, GOOD\}$, $\{LO, MED, HI\}$ respectively. The corresponding membership functions are considered to be triangular and can be mathematically expressed as

$$f_i(u, \phi_i) = \begin{cases} \frac{(u - \phi_{i1})}{(\phi_{i2} - \phi_{i1})} & \text{if } \phi_{i1} \leq u \leq \phi_{i2} \\ \frac{(u - \phi_{i3})}{(\phi_{i2} - \phi_{i3})} & \text{if } \phi_{i2} \leq u \leq \phi_{i3} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where, $u \in U_i$ is any value within the universe of discourse and $\phi_{ij} \in \phi_i \forall j$ are the parameters of the membership functions to be estimated by the MCMC algorithm. The rule base used for constructing the fuzzy inference system is as follows:

$$\begin{aligned} R_1 : & \text{if } loc_risk == HI \vee maintenance == POOR \\ & \text{then } downtime == HI \\ R_2 : & \text{if } loc_risk == MED \vee maintenance == AVG \\ & \text{then } downtime == MED \\ R_3 : & \text{if } loc_risk == LOW \wedge maintenance == GOOD \\ & \text{then } downtime == LOW \end{aligned} \quad (5)$$

For this case we assume all the rules to be known ($\beta_k = 1, \forall k$) and only the parameters (ϕ_i) of the membership functions (f_i) need to be estimated. For this example, 100 data points are generated from the rule base in Eq. (5) using triangular membership functions and centroid method of defuzzification. The values of the tips and the bases of the triangular membership functions are to be estimated by the MCMC algorithm. Therefore, for this case with three referential sets each with three referential attributes, there are 9 parameters to be estimated. Assume the following for the set of observations $\{\mathbf{Y}_N, \mathbf{X}_N\}$:

$$\mathbf{Y}_N \sim \mathcal{MVN}(g(\mathbf{X}_N; \theta), \sigma^2 I_N), \quad (6)$$

where $\mathcal{MVN}(\mu, \Sigma)$ is the Multivariate-Normal distribution with mean vector μ and covariance matrix Σ . The mean vector for our model comes from $g(\mathbf{X}_N; \theta)$ and the fuzzy inference system from Eq. (3). We use I_N , the $N \times N$ identity matrix, along with a parameter σ to construct the covariance matrix. In other words, the observations of 100 data-points are assumed to come from a normal distribution with the mean (μ) of the distribution given by the output of the fuzzy inference system

and a variance (σ) representing the measurement error. This leads to the likelihood

$$\begin{aligned} p(\mathbf{y}_N | g(\mathbf{x}; \theta), \mathbf{X}_N) \\ = \frac{1}{\sqrt{(2\pi)^N \sigma^{2N}}} \times \\ \exp\left(-\frac{1}{2\sigma^{2N}}(\mathbf{y}_N - g(\mathbf{X}_N; \theta))^T(\mathbf{y}_N - g(\mathbf{X}_N; \theta))\right). \end{aligned} \quad (7)$$

where \mathbf{y}_N is a realisation of the random variable \mathbf{Y}_N . The set of parameters is $\theta = \{\phi_i : i = 1 \dots n\}$, and we can estimate it by being Bayesian and assuming a prior $p(\theta)$. This gives the posterior distribution:

$$p(\theta | \mathbf{y}_N, \mathbf{X}_N) \propto p(\mathbf{y}_N | g(\mathbf{x}; \theta), \mathbf{X}_N) p(\theta). \quad (8)$$

Although Eq. (8) is not mathematically tractable, we can draw samples from it using a Metropolis-Hastings algorithm such as the Gibbs sampler [30].

a) *Case 1: Parameter estimation without measurement noise and small data set:* For this case, since the data-points are generated from the true fuzzy inference system without any addition of noise, we set $\sigma = 0.001$, i.e. the measurement error is negligible. We set the prior $p(\theta)$ to be independent uniform distributions over the parameter space for each ϕ_i

$$\phi_i \sim \text{Uniform}(0, C_{max}) \quad (9)$$

such that

$$\begin{aligned} p(\theta) &= p(\phi_1, \phi_2, \dots, \phi_n) \\ &= p(\phi_1)p(\phi_2) \dots p(\phi_n) \\ &= \left(\frac{1}{C_{max}}\right)^n. \end{aligned} \quad (10)$$

Here, C_{max} is the maximum allowed value for the parameter ϕ_i , see table IV for the values.

The first set of MCMC simulations are run with only 15 data points and the convergence diagnostics for one of the parameters (parameter 9) is shown below in Fig. 1. The convergence curves for all the others parameters are similar and the Highest Density Intervals (HDIs) are tabulated in Table IV in the appendix. For all the simulation cases, non-informative priors are used for all the parameters during the Bayesian estimation phase. For all the MCMC simulations in this paper, the python package pymc3 [31] is used.

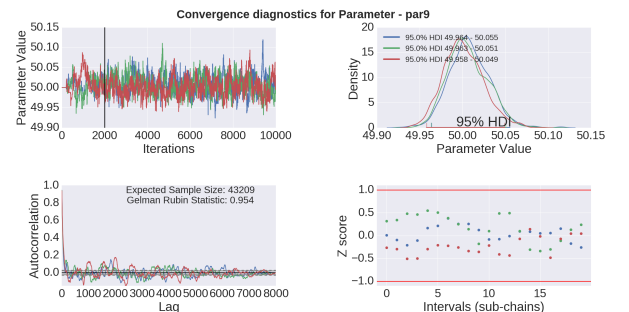


Fig. 1. Case I: MCMC convergence diagnostics for parameter 9 with 15 data points

The top left sub-plot in Fig. 1 shows the trace plots for three independent MCMC chains over the parameter space for 10,000 iterations. The burn-in is set at 2000 iterations and is indicated in the figure with a black line. The chains are well mixed and the density plots on the top right are overlapping and indicate that they are representative of the true posterior distribution. The 95% highest density intervals (HDI) are also plotted on the x-axis which represent the most probable range of the parameter value. For parameter 9, the true value was 50, and the HDIs for all the three chains are within 49.95 and 50.05, showing that the MCMC algorithm is able to correctly identify the true value of the parameter within a very tight range with only 15 data points. The convergence characteristics can also be checked with the auto-correlation plots of the three chains on the bottom left. The auto-correlation falls quickly with the number of lags for the post burn-in samples and therefore other methods like chain thinning are not required to obtain a good posterior distribution. The effective sample size (ESS) [15] is also reported in this sub-plot which essentially reflects the sample size of a completely non-autocorrelated chain which would yield the same information as the one obtained from the MCMC algorithm. A higher ESS indicates better approximation of the posterior distribution.

The Gelman-Rubin statistic or the shrink factor is a measure of the variance between the chains relative to the variance within the chains. The value for this case, as indicated in the figure, is near 1 which implies that the chains have converged. As a thumb rule, if the value is more than 1.1, then there are orphaned or stuck chains. The Z-scores given in the convergence diagnostic plots are the Geweke's diagnostic. We expect these to be from a normal distribution, so most values should be within the red lines in the range $[-1,1]$. If a substantial amount of the Z-scores fall outside of this, it implies a failure in convergence.

b) Case II: Parameter estimation without measurement noise and more data points: The next set of simulations are done with the same settings as the previous case but with 100 data-points instead of 15. The convergence characteristics of the same parameter (par9) is shown in Fig. 2 and as can be seen in the top right density sub-plot, the 95% HPD interval is very small around the true value of 50, demonstrating that more data-points help in obtaining much tighter estimates of the parameter values. The trace plots, auto-correlation plots and the Geweke scores indicate that the simulations have converged.

c) Case III: Parameter estimation with additional measurement noise: For this case, noise from a distribution of $\mathcal{N}(0,1)$ is added to the 100 observations. The likelihood function for the MCMC in Eq. (6) then has $\sigma = 1$. The convergence diagnostics with the density plots for the same parameter (par9) is shown in Fig. 3.

It can be seen that the HPD of the parameter estimate now spans a much larger region than the two previous simulation runs, due to the noisy data-points. The uncertainty in parameter estimates can be seen from the plots of the membership functions obtained by plugging in the values from the posterior distribution, as depicted in Fig. 4. It can be observed that

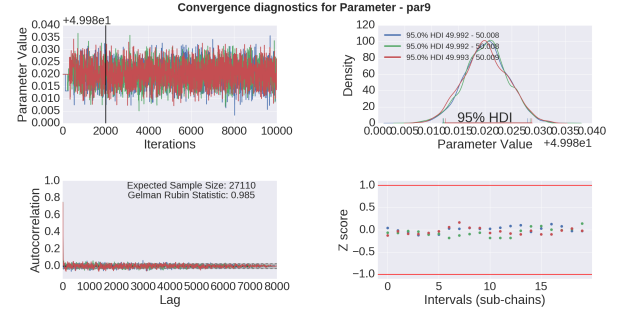


Fig. 2. Case II: MCMC convergence diagnostics for parameter 9 with 100 data points

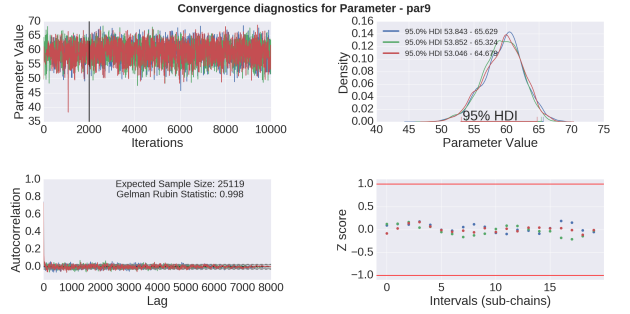


Fig. 3. Case III: MCMC convergence diagnostics for parameter 9 with 100 data points and added noise

for some of the membership function values like the LOW and HI membership function for Downtime, the uncertainties are much higher than the other ones. This might indicate that either there were not enough samples in that region to have more tighter estimates for these membership functions, or due to the non-linear nature of the fuzzy rule base, the output variable (Downtime) and its membership functions have more flexibility and can map the final output for a wider range of parameters.

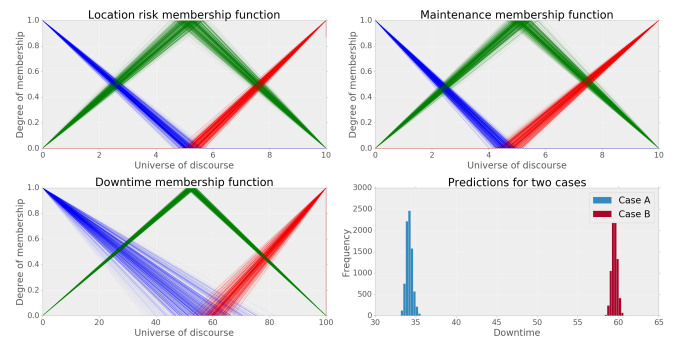


Fig. 4. Posterior of the membership functions and predictive distribution for two cases A and B. Fuzzy sets with blue colour indicate LO or POOR, green colour indicate MED or AVG and red colour indicate HI or GOOD.

The bottom right sub-plot in Fig. 4 shows the predictions of the Downtime for two different scenarios (Cases A and B). For Case A, the input variables are set to $[1,1,8,8]$ for location risk and maintenance respectively. For Case B, the values are

set to [7.7,1.1]. Therefore, in Case A the input suggests that the location risk is low while the maintenance is good. Hence we would intuitively expect the downtime to be low. For case B, the values suggest that the location risk is high and the maintenance is poor. Hence we would intuitively expect the Downtime to be high. This is confirmed by looking at the predictions obtained from the posterior samples as shown in the histograms at the bottom right sub-plot in Fig. 4. For Case A, the Downtime is low with a distribution around 34 and for Case B, the Downtime is high with a distribution around 59.

Once the parameter estimates are obtained, the prediction over the whole universe of discourse for all the variables can be obtained along with the posterior distribution of the uncertainty at each point. Fig. 5 shows the minimum and maximum surfaces of the predictions for the whole range of input parameters.




Fig. 5. The envelope of the minimum and maximum values of Downtime for Location Risk and Maintenance obtained by running the fuzzy function with the parameter samples from the posterior

2) *Bayesian rule base selection and parameter estimation of membership function (Case IV):* Consider the case where the dataset is generated as before from the three rules in Eq. (5), but two additional spurious rules are considered while estimating the parameters using MCMC. They are given by:

$$\begin{aligned} R_4 : & \text{if } loc_risk == LO \\ & \text{then } downtime == HI \\ R_5 : & \text{if } maintenance == POOR \\ & \text{then } downtime == LOW \end{aligned} \quad (11)$$

Clearly these two rules are contrary to our intuition of maintenance and breakdown events mapped on to system downtime and are antagonistic to the original three rules given by Eq. (5). For this case, we set up the MCMC to not only select the parameters of the membership functions (ϕ_i) but also the β_k in Eq. (3) which indicates the specific rules to be included for best explaining the dataset. Following the lines of [32], we set up a reversible jump MCMC by assuming a prior on the β_k parameters, which controls the inclusion of rule k . We assume that the inclusion of any rule is independent of the

inclusion of other rules, and that rules are equally likely to be included as excluded. Probabilistically, this is represented as:

$$\beta_k \sim \text{Bernoulli}(0.5) \quad (12)$$

and the β_k parameters can be sampled from the posterior distribution along with the other parameters.

As before, we ran the MCMC algorithm for 10,000 iterations and the convergence diagnostics of the same parameter (par9) is shown in Fig. 6. The effect of the burn-in is clear from the top left plot and the parameter estimates in the top right have a higher variance (compared to the previous cases) around the true value of 50. This is due to the fact that the same 100 data points are used to estimate the parameters as before, but in this case the number of parameters are greater – 14 parameters (9 parameters for the fuzzy membership functions and 5 parameters for selecting the individual rules) as opposed to just 9 parameters in the previous cases. The algorithm tries to extract more information from the data (more parameters to estimate) and hence the uncertainties in the parameters are higher. More data would decrease uncertainty around the estimates.

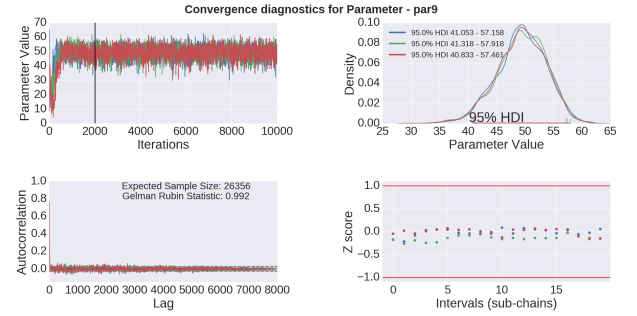


Fig. 6. Case IV: MCMC convergence diagnostics of parameter 9 with for the case of rule base selection

Since the rule selection parameter is a Bernoulli distribution (essentially 0 or 1 values instead of continuous real values), it is easier to interpret the results from a histogram plot (Fig. 7) of the MCMC samples for these parameters.




Fig. 7. Histogram of the rule base selection parameters. All values are binary (0 or 1), and due to the stacked nature of the plot they appear on the continuum of the abscissa.

Fig. 7 shows that the two spurious rules (Rule 4 & 5) were not chosen by the algorithm (most of the MCMC samples for these have zero values) and the other three rules which

generated the data (Rule 1, 2, 3) were correctly identified by the algorithm (most values are one). A small number of MCMC samples for Rules 4, 5 have values of one as well, which is during the initial exploration phase of the MCMC (the plot shows values for all the samples and not just post burn-in) and they quickly converge to the true values.

3) *Comparison of the method with Generalised Linear Models (GLM) using Bayesian parameter estimation:* To show the effectiveness of the proposed method, we compare four GLM structures with Bayesian parameter estimation. A GLM is composed of the following three elements: an error distribution, a link function, and a linear predictor [33]. We use the normal distribution as our error distribution:

$$Y_i \sim \mathcal{N}(\mu, \sigma^2) \quad (13)$$

We use the identity link function to link the expected value to the linear predictor. We test four GLMs, with linear predictors as given by the following equations, where $\alpha_i \forall i$ are the parameters to be estimated using MCMC.

$$GLM1: \mu = \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 \quad (14)$$

$$GLM2: \mu = \alpha_0 + \alpha_1 x_1 x_2 \quad (15)$$

$$GLM3: \mu = \alpha_0 + \alpha_1 x_1^2 + \alpha_2 x_2^2 + \alpha_3 x_1 x_2 \quad (16)$$

$$GLM4: \mu = \alpha_0 + \alpha_1 x_1 + \alpha_2 x_2 + \alpha_3 x_1^2 + \alpha_4 x_2^2 + \alpha_5 x_1 x_2 \quad (17)$$

For the MCMC sampling, the priors on the parameters were $\sigma \sim \text{half-Cauchy}(10)$ and $\alpha_i \sim \mathcal{N}(0, 20) \forall i$, which are noninformative priors, ensuring that most of the information would come from the data. The half-Cauchy distribution is defined as the Cauchy distribution left-truncated at zero, as given in [34]. The number of MCMC iterations are set to 2,000 and the burn-in as 500. The results are checked for convergence as before and the posterior samples after burn-in are plugged in to each of the GLMs and a distribution of the Downtime predictions over the whole 2D surface of Location Risk and Maintenance is obtained. The mean of these surfaces is plotted superimposed with the interpolated surface of the data points as shown in Fig. 8.

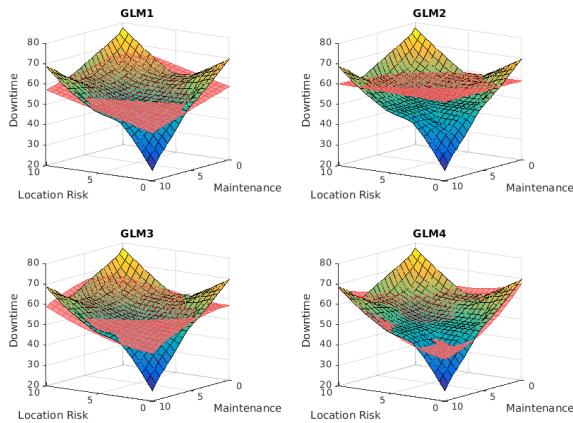


Fig. 8. Comparison of the mean surface of the predictions (red surface) from GLM posteriors with the interpolated surface of the actual data-points

As can be observed from Fig. 8, as more interactions and parameters are added in the GLM structure, it has more flexibility in fitting the data and consequently the fits are better, with GLM4 giving the closest fit among all others. However, as can be seen from the actual structure of GLM4 in Eq. 17, not much can be gleaned to relate to the expert opinions as expressed in Eq. 5 from which the data was generated. In essence, the GLM is one of the many proxies that can fit the data. Also since the surface plots are interpolated from the data-points themselves, they are an approximation of the real surface which generated the data and would be more closer to the real one as and when more data-points are gathered. This might also imply that as more data points are gathered, a different GLM structure might fit better. This can also happen with the fuzzy rules where the algorithm has to choose the rules (as shown in the previous section), but if the algorithm selects a different fuzzy rule on having more data points, there is a more transparent explanation of the dynamics in the model.

In short, it is difficult to come up with a GLM model structure to explain such non-linear relationship in the data. In general, modellers experiment with a combination of linear models, models with interactions, models with higher order terms etc. The only way to include expert opinion in a traditional GLM is to use priors on the parameters (which is a lot less flexible than the IF-ELSE structured fuzzy rule base method proposed in the paper). But since the structure of a GLM is decided ad-hoc it is unclear how to transform expert opinions into priors or model parameters.

B. Classification on synthetic data-sets

Our modelling methodology for regression can be easily extended for classification tasks. This extension is quite similar to using a logistic regression as a binary classifier. The binary classification task can be defined in terms of a Bernoulli likelihood function whose parameter (ψ) is set by the fuzzy inference system ($g(\cdot)$) as shown in Eq. 18.

$$Y_N \sim \text{Bernoulli}(\psi) \quad (18)$$

$$\psi = g(\mathbf{X}_N; \theta)$$

The rationale for this model for the classification task, is that the Bernoulli distribution gives either a one or a zero value (essentially representing the output class) depending on the probability parameter ψ . The fuzzy function acts like a non-linear kernel by taking in the input vectors \mathbf{X}_N and transforming it into a value in the range [0,1] using the pre-specified rule base and the corresponding parameters θ . The rule bases are specified as before (or can be learned using the rule base selection parameter) and the parameters of the fuzzy function can be estimated using the MCMC algorithm.

As an illustrative example, we use the same rule base as in Eq. 5 where the input co-variables are in the range [0-10], but the output is in the range [0,1] using triangular membership functions as in Eq. 4. We generate 100 data points from this fuzzy function and set any data point above a threshold of 0.5 to a value of 1 and any data point below it to a value of 0. This forms the training data-set for running the MCMC algorithm to

see if the true parameters θ of the fuzzy membership function can be recovered which would classify the data-points into the two classes. Fig. 9 shows the underlying surface generated with the true parameters θ , with superimposed 100 randomly drawn samples which is used as the training data-set.




Fig. 9. Underlying true surface from the fuzzy function which generates the two classes (red=1, blue=0)

For the simulations, 10,000 iterations are done for 3 independent chains and vague priors are chosen for θ as also done in all the regression cases. After rejecting 2000 samples for burn-in, the mean of the posterior parameter samples are plugged into the fuzzy function, with 0.5 as the threshold and the predicted class labels is plotted in Fig. 10.




Fig. 10. Binary classification results using the mean of the posterior samples in the Fuzzy Bayesian Learning approach. Predictions for Class 1 is in red, class 0 is in blue and the mis-classifications are shown with lime circles.

One advantage of this method is that the introduction of the fuzzy system acts as a non-linear kernel and helps in classifying data which are not necessarily linearly separable in the feature space. Also, class one (in red) has much more number of data points than the other (blue) and therefore this is

the so-called imbalanced data-set case [35]. The method is able to determine the classes without any special modification for imbalanced data-sets (e.g. changing objective function values to penalise mis-classification of one class more than the other etc). However, as explained earlier in the regression cases as well, the caveat with this method is that the rule bases might not always be trivial to find and choosing a totally wrong rule base would drastically affect the performance of the classifier.

Also, the class label prediction plots in Fig. 10 shows that there is a higher mis-classification at the class boundaries where the underlying fuzzy surface transits from greater than 0.5 to less than 0.5. This implies that the underlying fuzzy surface is not estimated properly with 100 data points, unlike the regression cases (with same or less number of data points) where the parameter estimates had really tight bounds. One possible reason might be that since the observations are just integer values (zeros or ones) instead of real numbers, there is less information for the MCMC algorithm to find the actual parameters with tighter uncertainty bounds.

C. Application to a real world problem

1) *Context and background of the problem:* The background of the problem is rooted in the financial services sector, specifically for making risk quantification models for the non-life speciality insurance domains. For this particular example, consider the case of insuring a power plant operation in the event of business interruption (BI) and property damage (PD) due to various risks like fire hazard, machinery breakdown etc. The focus of this paper is specifically the BI which depends on various factors like the reliability of each of the machinery components, redundancy schemes employed, maintenance schedules etc. Therefore, given a set of historic data on these co-variables and the corresponding system uptime in a year, it would be possible to fit a simple GLM or other regression models for prediction. The distribution of the system uptime can then be subsequently used in an insurance pricing model to obtain the required premium. This example focuses on better estimation of the system uptime for the insured engineering process.

It is possible to build more involved models of such engineering process systems in cases where individual failure and repair data of each of the underlying component is available as has been pursued in [36]. However, many plant operators are reluctant to share data and therefore such data might not be available for all different kind of process plants, to actually make a sizeable population so as to draw a statistical estimate of the failures and reliability of each component. The proposed Fuzzy Bayesian Learning (FBL) method is useful in such scenarios as it starts with some prior knowledge which informs the model parameters and some prior expert opinion which informs the model structure. As data comes in, the parameters can be re-estimated with these observations to get more tighter estimates around the model parameters and consequently predictions with less uncertainties.

This case also serves as a validation of the ability to use this kind of modelling for other type of emerging risk areas where no prior insurance product has been offered and consequently

there is little or no loss history data that can be used to build a data-driven model. In such cases, the fuzzy rule base relying on the underwriter's opinion without any data can be constructed to give a sense of the outputs. Later the Bayesian parameter estimation can be introduced as the actual claims data start coming in.

2) *Dataset description:* The data-set is collected in-house at Sciemus Ltd. London, for underwriting insurance risks for its power sector business. The data-set consists of 3 variables, which serve as an input to our model – operation & maintenance (O&M), loss history, design & technology & fire risk (DnTnF). These are qualitative opinions indicated by the underwriter from the submission document of the specific risk which comes in to get a quotation for insurance.

All these variables are on a scale of 0-10 with 0 representing very poor and 10 indicating very good. Since the dataset consists of historical information from the actual underwriting process, it is possible to see which risk actually led to a loss event and a consequent claim, what was the claim amount and the amount of system uptime during the given insurance period. For this problem setting, the actual outcome is the system uptime (percentage of time the system was working during the insurance period which is typically one year). This is a small data-set and Table III gives the summary statistics of the 24 data-points available in it.

3) *Simulations:* The fuzzy rule base used for simulation is shown in Table I. This is extracted from the intuition of the underwriter who believes that O&M rating is the key to system uptime. This can be seen reflected in the table – e.g. the first rule states that if the O&M is good, then irrespective of the loss history and design factors the system uptime would be high. The rationale behind such a rule is that if the plant engineers know how the plant operates and are very rigorous with the maintenance schedule and monitoring aspects then they might get a few failures now and then (and have some history of losses), but this is inherently a good risk to underwrite. In some cases though, the loss history can be the overriding factor, e.g. rule 4 states that if the loss history is good, then the underwriter gives precedence to this over the bad O&M. Similar lines of reasoning are extracted for the other rules as well. It is important to stress that multiple rules are triggered, each to a different extent, inside the fuzzy inference system, for one given antecedent input set. The final output is a combination of all these triggered values.

TABLE I
RULE BASE FOR THE FUZZY BAYESIAN MODEL FITTED ON THE REAL
WORLD CASE

Antecedents			Consequent
O&M	Loss history	DnTnF	System uptime
Good	-	-	High
Good	-	Average	Medium
Average	-	-	High
Bad	Good	-	High
Bad	Bad	Bad	Low

Nevertheless, this rule base extraction is not exhaustive and neither all the possible model spaces are explored to see if they

give similar results. A more rigorous approach to this would be Bayesian model comparison which involves comparing the evidence of different models to come up with the model that best explains the data. This is an investigation in its own right and a digression from the main theme of the paper which aims to introduce Fuzzy Bayesian Learning and therefore is not pursued in this paper. However, to show the effectiveness of this simple rule base, it is compared with the predictive accuracies of three GLMs (as shown in Eq. 21-23) fitted with MCMC techniques. The generic likelihood as specified in Eq. 13 is used for all the GLMs. The priors on each of the parameters are sufficiently non-informative to let the model learn from the dataset itself. The σ in Eq. 13 is set to a fixed value of 0.25 (as also done while estimating the parameters of the Fuzzy Bayesian model). Ideally this should have been estimated from the data-set itself, but estimating the variance with tight uncertainty bounds actually requires large number of data points and hence this is pre-specified to a small value. Doing so implies that we are relying on the linear predictor to explain most of the relationships in the underlying data-set.

For comparing among these models the mean squared error (MSE) is used to gauge the predictive accuracies. In Bayesian models, prediction of a new set of values $\tilde{\mathbf{y}}_N$ is done using the posterior predictive distribution:

$$\begin{aligned}
 p(\tilde{\mathbf{y}}_N | \mathbf{y}_N, \mathbf{X}_N) &= \int p(\tilde{\mathbf{y}}_N, \theta | \mathbf{y}_N, \mathbf{X}_N) d\theta \\
 &= \int p(\tilde{\mathbf{y}}_N | \theta, \mathbf{y}_N, \mathbf{X}_N) p(\theta | \mathbf{y}_N, \mathbf{X}_N) d\theta \\
 &= \int p(\tilde{\mathbf{y}}_N | \theta, \mathbf{X}_N) p(\theta | \mathbf{y}_N, \mathbf{X}_N) d\theta.
 \end{aligned} \tag{19}$$

Here, $p(\tilde{\mathbf{y}}_N | \theta, \mathbf{X}_N)$ is the likelihood assumed on the model in Eq. (8), and $p(\theta | \mathbf{y}_N, \mathbf{X}_N)$ is the posterior distribution. Drawing samples from the posterior predictive distribution is simple when we already have draws from the posterior distribution as provided by the MCMC sampler [15]. We used the mean of the posterior predictive distribution, $\mathbf{y}_N^{\text{pred}} = \mathbb{E}[\tilde{\mathbf{y}}_N]$ as our final prediction and compare it with the real data-point by calculating the mean squared error (MSE):

$$\text{MSE} = \frac{1}{N} \left(\mathbf{y}_N^{\text{pred}} - \mathbf{y}_N \right)^T \left(\mathbf{y}_N^{\text{pred}} - \mathbf{y}_N \right) \tag{20}$$

However, the MSE metric being a single number does not convey the uncertainties in the predictions properly, which can be gleaned from the whole posterior distribution of the prediction (but this again is difficult to compare with each other).

The data-set is intentionally not divided into a training and a test set to perform hold out cross validation or k-fold cross validation, since this is a small data-set and ideally we would like the model not to lose out on much information so as to estimate the parameters with more certainty.

$$\text{GLM5} : \mu = \sum_i \alpha_i x_i \tag{21}$$

$$\text{GLM6} : \mu = \sum_i \alpha_i x_i + \sum_{\substack{i,j \\ i \neq j}} \alpha_{ij} x_i x_j \tag{22}$$

$$GLM7: \mu = \sum_i \alpha_i x_i + \sum_{\substack{i,j \\ i \neq j}} \alpha_{ij} x_i x_j + \sum_{\substack{i,j,k \\ i \neq j \neq k}} \alpha_{ijk} x_i x_j x_k \quad (23)$$

The resulting comparative analysis (using MSE as a metric), between the different models on the real world problem is shown in Table II. As can be observed, the Fuzzy-Bayesian model has a better MSE than GLM 5 which has a simple structure, but is worse off as compared to the other two GLMs (6 and 7). However, as discussed before, there is a narrative and more open interpretation for why the Fuzzy-Bayesian model works as opposed to the GLMs. Also predictive accuracy from such a small data-set is not really indicative of the performance of the model on the true underlying process and obtaining a lower MSE on this data-set might not always be the best option in the long run. As already discussed before, with new data points, different GLMs might give lower MSEs. However, if the underlying process can be quantified with the fuzzy rule bases, then this would lead to a good fit in subsequent scenarios.

TABLE II
COMPARISON OF MEAN SQUARED PREDICTION ERRORS FOR THE FUZZY BAYESIAN MODEL VIS-A-VIS OTHER GLMS

Model	MSE
GLM 5	0.4098
GLM 6	0.3649
GLM 7	0.3596
Fuzzy Bayesian	0.3654

The advantage of FBL in the real world example is that the underwriting insights can make the model better. But the other downside is that if the rules are extracted only from one underwriter, then their own biases and risk perceptions might adversely affect the model. To alleviate such concerns, the rule base can be constructed by consulting multiple underwriters and as more data comes in, the individual rule selection methodology (as shown in Section II-A2) can be applied in the parameter estimation process to reject the rules that do not describe the data properly. This allows the model to leverage on the expertise of the best underwriters.

III. DISCUSSIONS AND CHALLENGES

A. Relationship with other machine learning methods

Most machine learning methods cannot intuitively explain the fitted model in terms of the underlying assumptions. For example, a neural network might give very good prediction accuracy on a data-set, but not much reasoning can be gleaned from the model, just by looking at the learned weights of the individual neurons. This might work in applications in an automated framework (like recommender systems, user behaviour prediction etc.) where large number of decisions need to be made in very short time scales without human interventions, and the financial liabilities associated with a mis-classification is not that high. However, in other industries like personalised healthcare/medicine or specifically our real world example of speciality non-life insurance, prediction from a computer model is almost always validated by a human before taking

a final decision, since the consequences of such decisions have huge financial/legal implications. In such circumstances, having a deeper insight into why the algorithm actually works (as opposed to a black-box which just spits out gospel) makes it easier to take a decision and is always preferable. This issue is addressed by our proposed framework since the fuzzy rule base derived from the experts is used to inform the model structure for the regression or classification task. Also the associated uncertainties in the predictions would help the human decision maker to either ask for more information (i.e. recommend to do more medical tests for the healthcare case) or arrive at a different decision altogether due to the higher associated uncertainties (e.g. decide not to insure a specific risk or charge higher premiums to reflect the higher undertaking of risk).

The other issue with traditional machine learning methods, is that they require quite significant amount of data to train and cross-validate a model properly. For the real world cases described above, even though the expert knows the dependence between the co-variates, it is almost always difficult and expensive to collect a comprehensive data-set and therefore modelling methodologies have to work with small data-sets and still give reasonable predictions. These kind of application areas would also be the ones where the proposed methodology would work well. The advantage would also be that we can start with just the opinions of the expert without any data (there would be more uncertainties in the predictions) and then update the model with the data as we get more information (uncertainties in the predictions would decrease). The use of Fuzzy Bayesian Learning allows for such flexibility.

B. Effect of informative priors

In all the simulations, we have used noninformative priors. It might be especially useful to incorporate informative priors when there is even less data-points (e.g. underwriting emerging risks in a new sector which has only been done a few times or not at all). Judicious choice of the informative priors would help in giving smaller variance around the predictions. The predictive outputs for a Bayesian case is always a trade-off between the priors and the observed data. Therefore, as more data comes in, the model outputs are dictated by the data and the effect of these informative priors would diminish.

C. Improvement over traditional GLMs

This work can be seen as an extension of traditional GLMs with Bayesian parameter estimation. The proposed methodology retains all the advantage of Bayesian GLMs (e.g. simple model, intuitive explanation of parameters, inclusion of prior knowledge in the form of probability distribution), but makes the modelling methodology more flexible in the sense that the arbitrary choice of link functions can now be replaced with a model structure from the judgement of experts. Also qualitative imprecise opinions from the experts can be used to drive the model. Therefore, the model can handle both imprecision and uncertainty in the same framework.

In fact, the general model structure is the same for our proposed model and a GLM. Let y be a response variable,

and let \mathbf{x} be a set of covariates observed alongside it. Let μ and σ^2 be the mean and variance of a chosen error distribution, respectively. We construct the general model structure as:

$$\begin{aligned} y &\sim \text{Error-Distribution}(\mu, \sigma^2) \\ \mu &= g(\mathbf{x}; \boldsymbol{\theta}) \end{aligned} \quad (24)$$

where $\boldsymbol{\theta}$ is the set of parameters to be estimated in the model. For a GLM with k covariates, we have:

$$g(\mathbf{x}; \boldsymbol{\theta}) = h\left(\theta_0 + \sum_{i=1}^k \theta_i x_i\right) \quad (25)$$

where $h(\cdot)$ is known as the inverse-link function, and it connects a linear combination of the parameters to the mean of the response variable. In this paper, $g(\mathbf{x}; \boldsymbol{\theta})$ is the fuzzy system defined by the rulebase in Eq. (3), so we replaced the inverse-link function and the linear predictor in the GLM.

D. Relationship with type-2 and higher fuzzy sets

A type-2 fuzzy set can incorporate uncertainty about its membership functions using a third dimension and its associated footprint of uncertainty (FOU). For interval type-2 fuzzy sets, the third dimension has the same value everywhere and can be completely defined by an upper & lower membership function (UMF and LMF respectively) both of which are fuzzy type-1 sets. For this case, the area between the UMF and the LMF essentially represents the uncertainty. The Bayesian method can be used to learn from the data and reduce the area representing the uncertainties, thereby improving the predictions.

This framework therefore complements higher order fuzzy sets in the sense that irrespective of how the uncertainty is defined, the Bayesian method would act as a wrapper on top of the fuzzy logic based model, to reduce the band of uncertainty by learning from the data as more data-points come in. In terms of the present framework, defining an uncertainty range for the membership functions is somewhat similar to setting informative priors, instead of the uniform priors that we have considered in our simulations.

The other key difference is that the fuzzy type-2 method uses a type reducer between the inference and the defuzzification steps and the final output is always crisp, even though the membership functions have uncertainties in them. However, the proposed framework gives a posterior probability distribution for the predictions which helps in understanding the uncertainties associated with the predictions and aids in better decision making. It can also learn from new data that comes in and decrease the uncertainties in the membership functions and consequently in the predictions.

E. Potential problems and unresolved issues with the methodology

Since the model structure is shaped by the rule base provided by the experts, the predictive accuracy would be dependent on how good these rules are at explaining the data. Therefore, this methodology would work well in cases with small to medium data sets where the expert has good

understanding of the problem domain. For other cases where there is a large number of data and there is not much domain knowledge, traditional machine learning methods like neural networks, support vector machines etc. would be a better alternative.

Proper choice of the specific rules would also affect the uncertainties in the parameter estimates and consequently in the predictions. Each rule would add an extra rule selection parameter (β_k) to be estimated by the MCMC from the data-set. If the data-set is small, estimation of more β_k 's indicate that the uncertainty in the parameter estimates would be higher. On the other hand if the rules are not exhaustive enough to explain the data, then the predictive accuracy of the model would be poor as the search space of the MCMC is effectively limited by this rule base set. Therefore, succinct and judicious choice of rules are important for smaller predictive uncertainties and better predictive accuracies.

Even though the present work is geared primarily towards small and medium data-sets, the same methodology might be useful for large data-sets which do not fit in memory (i.e. big data). The rationale for this might be along similar lines to that of the present paper, where the interpretability of the model is also important vis-à-vis the predictive accuracy. However, Bayesian methods & particularly MCMC is inherently a sequential algorithm and hence would not directly scale to huge datasets. However, partitioning the data and training the model on a subset of the data and recursively estimating the parameters by doing block updates might be a solution, i.e. the posterior distribution of the parameters estimated from one subset of the data becomes the prior distribution for the next subset. Other sampling methods or variational inference might also help in speeding up the computations, but more research needs to be done in this area.

Including the fuzzy inference system and calling it at each MCMC iteration inherently makes the computation times much higher. Other solutions like developing surrogates for the expensive function and sampling from it only when required might be an avenue to explore. However, this might introduce additional errors in the parameter estimates as the surrogate model itself would be an approximation of the true function. Further research needs to be done on this.

F. Directions for future research

Future improvisations can look at recursive updating and sequential monte carlo methods (like particle filtering, unscented Kalman filtering etc.) so that when a new data-point arrives, the algorithm does not need to run again from scratch with the whole data-set and just updates the parameter values from the posterior of the old data-set and the new observation value.

Another direction of research can be to look at more efficient sampling techniques than the present Metropolis-Hastings one, to improve convergence characteristics. Gradient free samplers like Slice sampling can be used for this case, but most other improved sampling like Hamiltonian Monte Carlo etc. would require gradients from the black box fuzzy function, which is difficult to obtain.

Improvisations could also be made in the direction of computational intelligence, like extending the methodology

to other kinds of fuzzy sets like intuitionistic fuzzy sets, non-stationary fuzzy sets, type-2 fuzzy sets, hesitant fuzzy sets etc. The fuzzy rule bases could be evolved through genetic programming or some variant of genetic algorithm [37], however this would potentially require more data for the learning process or might result in overfitting and consequent loss of generalisation capabilities of the model. Automated rule generation processes might also result in logically incoherent rule bases (even though they might have high predictive accuracy) and also loss of coherent reasoning behind the choice of such rules [38].

IV. CONCLUSION

This paper introduced the symbiotic confluence of Fuzzy logic and Bayesian inference to learn from data and also incorporate the model structure in the form of an expert rule base. This helps in providing a narrative and transparency regarding the workings of the model while being able to include complicated non-linear transformations (in the form of rule bases) which give good predictive accuracy on a data-set. Unlike the generic fuzzy logic models, the outputs from the Fuzzy-Bayesian models are probabilistic and as demonstrated in the paper, can be used for regression, classification and also rule base selection tasks. Apart from synthetic examples, the approach is also demonstrated on a real life application in the financial services sector and is especially suited to small and medium data-sets. Extensive discussions are included for the relationship of this methodology with other machine learning methods and potential pitfalls like improper choice of rule-bases which would hamper the predictive capabilities of the model. Future directions for research involve using higher order fuzzy sets, improved sampling algorithms, recursive estimation and surrogate model based techniques for computational run-time reduction.

ACKNOWLEDGMENT

The authors would like to thank Sciemus Ltd. for sponsoring this work. They would also like to thank Angelo Nicotra for helping with the underwriting information gathering process and Stefan Geisse for his insights in the power underwriting business.

REFERENCES

- [1] L. A. Zadeh, "Discussion: Probability theory and fuzzy logic are complementary rather than competitive," *Technometrics*, vol. 37, no. 3, pp. 271–276, 1995.
- [2] R. R. Yager, "A note on probabilities of fuzzy events," *Information Sciences*, vol. 18, no. 2, pp. 113 – 129, 1979.
- [3] L. A. Zadeh, "Special issue information theory applications to problems of information science fuzzy probabilities," *Information Processing & Management*, vol. 20, no. 3, pp. 363 – 372, 1984.
- [4] —, "A note on z-numbers," *Information Sciences*, vol. 181, no. 14, pp. 2923 – 2932, 2011.
- [5] D. Dubois and H. Prade, *Fuzzy sets and systems: theory and applications*. Academic Press, Access Online via Elsevier, 1980.
- [6] M. Mizumoto and T. K., "Some properties of fuzzy sets of type 2," *Information and Control*, vol. 31, pp. 312–340, 1976.
- [7] K. Atanassov, "Intuitionistic fuzzy sets," *Fuzzy Sets and Systems*, vol. 20, pp. 87–96, 1986.
- [8] J. M. Garibaldi, M. Jaroszewski, and S. Musikasuan, "Nonstationary fuzzy sets," *IEEE Transactions on Fuzzy Systems*, vol. 16, no. 4, pp. 1072–1086, 2008.
- [9] R. R. Yager, "On the theory of bags," *International Journal of General System*, vol. 13, no. 1, pp. 23–37, 1986.
- [10] V. Torra, "Hesitant fuzzy sets," *International Journal of Intelligent Systems*, vol. 25, no. 6, pp. 529–539, 2010.
- [11] L. A. Zadeh, "The concept of a linguistic variable and its application to approximate reasoning," *Information sciences*, vol. 8, no. 3, pp. 199–249, 1975.
- [12] Y. Dong, Y. Xu, and S. Yu, "Computing the numerical scale of the linguistic term set for the 2-tuple fuzzy linguistic representation model," *IEEE Transactions on Fuzzy Systems*, vol. 17, no. 6, pp. 1366–1378, 2009.
- [13] F. Herrera and L. Martínez, "A 2-tuple fuzzy linguistic representation model for computing with words," *IEEE Transactions on fuzzy systems*, vol. 8, no. 6, pp. 746–752, 2000.
- [14] R. M. Rodríguez, L. Martínez, and F. Herrera, "Hesitant fuzzy linguistic term sets for decision making," *IEEE Transactions on Fuzzy Systems*, vol. 20, no. 1, pp. 109–119, 2012.
- [15] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian data analysis*, 3rd ed. Taylor & Francis, 2014.
- [16] Y. Li, J. Chen, and L. Feng, "Dealing with uncertainty: a survey of theories and practices," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 11, pp. 2463–2482, 2013.
- [17] M. M. Rajabi and B. Ataie-Ashtiani, "Efficient fuzzy bayesian inference algorithms for incorporating expert knowledge in parameter estimation," *Journal of Hydrology*, vol. 536, pp. 255–272, 2016.
- [18] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi, "Multisensor data fusion: A review of the state-of-the-art," *Information Fusion*, vol. 14, no. 1, pp. 28–44, 2013.
- [19] J. L. Ross, M. M. Ozbek, and G. F. Pinder, "Aleatoric and epistemic uncertainty in groundwater flow and transport simulation," *Water Resources Research*, vol. 45, no. 12, 2009.
- [20] R. Viertl, *Is it Necessary to Develop a Fuzzy Bayesian Inference ?* Boston, MA: Springer US, 1987, pp. 471–475. [Online]. Available: http://dx.doi.org/10.1007/978-1-4613-1885-9_48
- [21] —, *Fuzzy Bayesian Inference*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 10–15. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-85027-4_2
- [22] S. Frühwirth-Schnatter, "On fuzzy bayesian inference," *Fuzzy Sets and Systems*, vol. 60, no. 1, pp. 41–58, 1993.
- [23] M. Arefi and S. M. Taheri, "Possibilistic bayesian inference based on fuzzy data," *International Journal of Machine Learning and Cybernetics*, pp. 1–11, 2010.
- [24] J.-B. Yang, J. Liu, J. Wang, H.-S. Sii, and H.-W. Wang, "Belief rule-base inference methodology using the evidential reasoning approach-rimer," *IEEE Transactions on systems, Man, and Cybernetics-part A: Systems and Humans*, vol. 36, no. 2, pp. 266–285, 2006.
- [25] J.-B. Yang, J. Liu, D.-L. Xu, J. Wang, and H. Wang, "Optimization models for training belief-rule-based systems," *IEEE Transactions on systems, Man, and Cybernetics-part A: Systems and Humans*, vol. 37, no. 4, pp. 569–585, 2007.
- [26] Z.-J. Zhou, C.-H. Hu, D.-L. Xu, J.-B. Yang, and D.-H. Zhou, "Bayesian reasoning approach based recursive algorithm for online updating belief rule based expert system of pipeline leak detection," *Expert Systems with Applications*, vol. 38, no. 4, pp. 3937–3943, 2011.
- [27] O. Osoba, S. Mitaim, and B. Kosko, "Bayesian inference with adaptive fuzzy priors and likelihoods," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 41, no. 5, pp. 1183–1197, 2011.
- [28] W. Gilks, S. Richardson, and D. J. Spiegelhalter, "Introducing Markov chain Monte Carlo," in *Markov Chain Monte Carlo in Practice*, W. Gilks, S. Richardson, and D. Spiegelhalter, Eds. Chapman & Hall/CRC, 1996, ch. 1, pp. 1–19.
- [29] D. J. MacKay, *Information theory, inference and learning algorithms*. Cambridge university press, 2003.
- [30] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. Springer-Verlag New York Inc., 2004.
- [31] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, "Probabilistic programming in python using pymc3," *PeerJ Computer Science*, vol. 2, p. e55, 2016.
- [32] D. J. Lunn, N. Best, and J. C. Whittaker, "Generic reversible jump MCMC using graphical models," *Statistics and Computing*, vol. 19, no. 4, pp. 395–408, 2009.
- [33] P. McCullagh and J. Nelder, *Generalized Linear Models, Second Edition*, ser. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis, 1989.
- [34] A. Gelman *et al.*, "Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper)," *Bayesian analysis*, vol. 1, no. 3, pp. 515–534, 2006.

- [35] H. He and E. A. Garcia, "Learning from imbalanced data," *IEEE Transactions on knowledge and data engineering*, vol. 21, no. 9, pp. 1263–1284, 2009.
- [36] I. Pan, A. Korre, and S. Durucan, "A systems based approach for financial risk modelling and optimisation of the mineral processing and metal production industry," *Computers & Chemical Engineering*, vol. 89, pp. 84–105, 2016.
- [37] O. Cordon, F. Herrera, and P. Villar, "Generating the knowledge base of a fuzzy rule-based system by the genetic learning of the data base," *IEEE Transactions on fuzzy systems*, vol. 9, no. 4, pp. 667–674, 2001.
- [38] J. Casillas, O. Cordon, F. H. Triguero, and L. Magdalena, *Interpretability issues in fuzzy modeling*. Springer, 2013, vol. 128.

APPENDIX A

TABLE III
SUMMARY STATISTICS FOR INPUT DATA

	O&M	Loss history	DnTnF	System uptime
count	24	24	24	24
mean	6.71	6.42	5.71	8.54
std	1.65	1.93	1.49	2.52
min	4.00	2.00	2.00	2.00
25%	6.00	5.00	5.00	7.75
50%	7.00	6.50	6.00	10.00
75%	8.00	7.00	6.25	10.00
max	9.00	10.00	8.00	10.00

Cases for Tables IV and V:

Case I: Regression with small dataset and no noise

Case II: Regression with bigger dataset and no noise

Case III: Regression with noise

Case IV: Regression with rule base selection

Note: Case IV also has the rule selection parameters which are binary and hence HDIs are not reported

TABLE IV
HIGHEST POSTERIOR DENSITY ESTIMATES (L= LOWER RANGE OF HDI, H=HIGHER RANGE OF HDI) FOR REGRESSION CASES I-IV

		Case I		Case II		Case III		Case IV	
		HDI (L)	HDI (H)	HDI (L)	HDI (H)	HDI (L)	HDI (H)	HDI (L)	HDI (H)
Param 1	Chain 1	4.997	5.002	5.000	5.000	4.761	5.620	4.587	5.476
	Chain 2	4.997	5.003	5.000	5.000	4.768	5.600	4.592	5.468
	Chain 3	4.997	5.002	5.000	5.000	4.779	5.602	4.590	5.501
Param 2	Chain 1	4.998	5.003	4.999	5.000	4.650	5.676	4.528	5.482
	Chain 2	4.998	5.003	5.000	5.000	4.657	5.686	4.469	5.502
	Chain 3	4.998	5.003	5.000	5.000	4.651	5.654	4.475	5.478
Param 3	Chain 1	4.999	5.001	5.000	5.000	4.948	5.661	4.608	5.303
	Chain 2	4.999	5.001	5.000	5.000	4.906	5.614	4.641	5.369
	Chain 3	4.999	5.001	5.000	5.000	4.947	5.649	4.650	5.336
Param 4	Chain 1	4.999	5.001	5.000	5.000	4.188	4.964	4.603	5.480
	Chain 2	4.999	5.001	5.000	5.000	4.198	5.002	4.553	5.415
	Chain 3	4.999	5.001	5.000	5.000	4.209	5.051	4.510	5.417
Param 5	Chain 1	4.997	5.002	4.999	5.000	4.682	5.523	4.573	5.412
	Chain 2	4.998	5.002	5.000	5.000	4.650	5.543	4.539	5.389
	Chain 3	4.998	5.002	5.000	5.000	4.699	5.569	4.578	5.432
Param 6	Chain 1	4.999	5.001	4.999	5.001	4.205	5.222	4.509	5.485
	Chain 2	4.999	5.001	4.999	5.001	4.205	5.241	4.481	5.442
	Chain 3	4.999	5.001	4.999	5.001	4.250	5.222	4.452	5.448
Param 7	Chain 1	49.922	50.060	49.970	50.031	41.959	72.715	35.040	83.099
	Chain 2	49.927	50.063	49.975	50.034	39.990	71.748	33.714	82.078
	Chain 3	49.931	50.068	49.970	50.030	40.781	73.217	33.257	81.149
Param 8	Chain 1	49.995	50.006	49.998	50.002	50.788	54.042	48.119	51.368
	Chain 2	49.995	50.006	49.999	50.002	50.785	53.874	48.499	51.657
	Chain 3	49.994	50.006	49.998	50.002	50.634	53.941	48.540	51.846
Param 9	Chain 1	49.964	50.055	49.992	50.008	53.843	65.629	41.053	57.158
	Chain 2	49.963	50.051	49.992	50.008	53.852	65.324	41.318	57.918
	Chain 3	49.958	50.049	49.993	50.009	53.046	64.679	40.833	57.461

TABLE V
NUMERICAL CONVERGENCE DIAGNOSTICS FOR THE REGRESSION CASES I-IV, (ESS = EFFECTIVE SAMPLE SIZE, GL STAT = GELMAN RUBIN STATISTIC)

	Case I		Case II		Case III		Case IV	
	ESS	GL stat	ESS	GL stat	ESS	GL stat	ESS	GL stat
Param 1	26256	0.986	24922	0.995	22599	1.006	25721	0.995
Param 2	35320	0.964	24960	1.000	23192	0.997	23239	0.999
Param 3	21440	1.019	26320	0.990	24140	1.002	26075	0.990
Param 4	24326	0.997	26535	0.991	22565	1.007	27549	0.990
Param 5	40145	0.959	24867	0.998	24191	0.993	25653	0.995
Param 6	26620	0.992	25598	0.989	24643	0.996	24795	0.998
Param 7	41571	0.958	22331	1.012	23611	1.002	25420	0.991
Param 8	25070	0.995	27746	0.984	23939	1.004	27761	0.992
Param 9	43209	0.954	27110	0.985	25119	0.998	26356	0.992